

REMARKS

Applicant has received the Notice of Panel Decision dated December 11, 2006. This Preliminary Amendment is in response to the Notice of Panel Decision and the Office action dated September 22, 2006 in which the Examiner: 1) rejected claim 1 as allegedly failing to comply with the written description requirement; 2) rejected claim 1 as allegedly indefinite; 3) rejected claim 1 as allegedly directed to non-statutory subject matter; 4) rejected claims 15-18 as allegedly anticipated by Mikesell (U.S. Pub. No. 2004/0153479, hereinafter "Mikesell"); 5) rejected claims 8-14 and 19-25 as allegedly unpatentable over Cannon (U.S. Pat. No. 5,983,239, hereinafter "Cannon") and Mikesell; and 6) rejected claims 1-7 as allegedly unpatentable over Cannon, Howard (U.S. Pat. No. 6,519,612, hereinafter "Howard") and Mikesell.

With this Response, Applicant amends claims 1-14 and 16-18. Reconsideration is respectfully requested.

I. WRITTEN DESCRIPTION REJECTION

The Notice of Panel Decision indicates the withdrawal of the written description rejection of Claim 1.

II. INDEFINITENESS REJECTION

With this Preliminary Amendment, Applicant amends claim 1 to address both aspects of the indefiniteness rejection.

III. NON-STATUTORY SUBJECT MATTER REJECTION

The Office action of September 22, 2006 rejects claim 1 as allegedly being directed to non-statutory subject matter. According to the MPEP, descriptive material comes in two varieties: functional descriptive material; and non-functional descriptive material. (MPEP 8th Ed. Rev. 3, August 2005, § 2106(IV)(B)(1), pp. 2100-11, 12).

[F]unctional descriptive material' consists of data structures and **computer programs** which impart functionality when employed as a computer component.

(*Id.* (emphasis added)).

When functional descriptive material is recorded on some computer-readable medium it becomes structurally and functionally interrelated to the medium and will be statutory in most cases since use of technology permits the function of the descriptive material to be realized.

(*Id.*(emphasis added)).

Claims 1-7 are now “computer-readable medium” claims. In order to avoid any assertion that claims 1-7 are still method claims, Applicant amends claim 1-7 to remove the method-like terminology. Thus, claims 1-7 are directed to a physical object (computer-readable medium) storing functional descriptive material in the form of a computer program. For this reason the Section 101 rejection should be withdrawn.

IV. ART-BASED REJECTIONS

A. Claim 1

Claim 1 stands rejected as allegedly obvious over Cannon and Howard. Applicant amends claim 1 to more fully address the allegedly non-statutory nature of the claim, and not to define over the cited art.

Cannon is directed to a storage management system with file aggregation supporting multiple aggregated file counterparts. (Cannon Title). In particular, Cannon discusses the existence of processing overhead associated with a file in a system, and that for smaller files the processing overhead is a predominant factor in accessing the smaller files.

The storage of each file requires both media preparation overhead and bookkeeping overhead, delaying completion of the entire storage process. The overhead for storage of a file is independent of that file's size. Thus, the overhead for a large file is overshadowed by its more substantial I/O time. The opposite is true with small files, where the necessary overhead dominates the file storage process compared to the file's relatively short I/O time. Consequently, I/O time is the chief obstacle in speedier storage of large files, whereas overhead prevents small files from being stored faster.

(Cannon Col. 1, line 65 through Col. 2, line 7). In order to decrease the effect of the processing overhead associated with each access to a file, Cannon discloses

a data storage subsystem 102 that performs services for client stations 106 (e.g., archive and backup) where user files are aggregated into managed files.

One of the key features of the present invention is storage and use of **"managed" files, each comprising an aggregation of one or multiple constituent "user" files.** The "user" files are created by the client stations 106, and managed by the subsystem 102 as a service to the client stations 106. ... This "internal" management scheme [of the subsystem 102] helps to significantly reduce file management overhead costs by using managed files constructed as aggregations of many different user files. In particular, **the subsystem 102 treats each managed file as a single file during backup, move, and other subsystem operations, reducing the file management overhead to that of a single file.**

(Cannon Col. 7, lines 54-67 (emphasis added)). Cannon discloses tables that track the membership of various user files in a managed file; however, the user files are consistently identified by the name assigned by client stations 106. (Cannon Col. 8, lines 29 through Col. 9, line 11; Table 1 (note "user file name" column); Table 3 (note "user file name" column)).

Claim 1, by contrast, specifically recites a computer-readable medium storing a program that "receive[s] a file from a client machine by appearing to operate in a client machine namespace and in a client machine file structure; ... and implement[s], autonomously of a user of the file, storage strategies for the file based on the metadata and in a namespace different than the client machine namespace." Applicant respectfully submits that Cannon and Howard do not teach or suggest such a system. In particular, Cannon appears to teach tracking a file throughout the Cannon system by the file name assigned by the client station. Howard is cited only for "implementing storage strategies autonomously of the user." (Office action of April 7, 2006, page 9, numbered paragraph 9). Thus, even if hypothetically the teachings of Howard are precisely as the Office actions suggests (which Applicant does not admit is proper), Cannon and Howard still fail to teach or suggest programs that "implement, autonomously of a user of the file, storage strategies for the file based on the metadata and in a namespace different than the client machine namespace."

Based on the forgoing, Applicant respectfully submits that claim 1, and all claims which depend from claim 1 (claims 2-7), should be allowed. Applicant amends claims 2-7 to reflect the amendment regarding the alleged non-statutory subject matter, and not to define over any cited art.

B. Claim 8

Claim 8 stands rejected as allegedly obvious over Cannon and Mikesell. Applicant amends claim 8 to remove the “adapted to” terminology and to remove reference to the program to make the claim more clear, and not to define over the cited art.

Cannon is directed to a storage management system with file aggregation supporting multiple aggregated file counterparts. (Cannon Title). In particular, Cannon discusses the existence of processing overhead associated with a file in a system, and that for smaller files the processing overhead is a predominant factor in accessing the smaller files. (Cannon Col. 1, line 65 through Col. 2, line 7). In order to decrease the effect of the processing overhead associated with each access to a file, Cannon discloses a data storage subsystem 102 that performs services for client stations 106 (*e.g.*, archive and backup) where user files are aggregated into managed files. (Cannon Col. 7, lines 54-67). Cannon discloses tables that track the membership of various user files in a managed file; however, the user files appear to consistently be identified by the name assigned by client stations 106. (Cannon Col. 8, lines 29 through Col. 9, line 11; Table 1 (note “user file name” column); Table 3 (note “user file name” column)).

Claim 8, by contrast, specifically recites, “wherein the host computer communicates files to the server for storage on at least one of the plurality of storage devices, wherein the server appears to be a network storage device operating in a user name space and in a user file structure; and wherein the server selects on which of the plurality of storage devices to store the files on a file-by-file basis based on storage characteristic preferences supplied for each file, and wherein each file is stored under a globally unique name in a global namespace of the server.” Applicant respectfully submits that Cannon and Mikesell do not teach or suggest such a system. In particular, Cannon appears to

teach tracking a file throughout the Cannon system by the file name assigned by the client station. Thus, even if hypothetically the teachings of Mikesell are precisely as the Office action suggests (which Applicant does not admit), Cannon and Mikesell fail teach or suggest “and wherein each file is stored under a globally unique name in a global namespace of the server.”

Based on the forgoing, Applicant respectfully submits that claim 8, and all claims which depend from claim 8 (claims 9-14), should be allowed. Applicant amends claims 9-14 consistent with the amendments to claim 8, and not to define over the cited art.

C. Claim 15

Claim 15 stands rejected as allegedly anticipated by Mikesell.

Mikesell is directed to systems and methods for restriping files in a distributed file system. (Mikesell Title). In particular, Mikesell is directed to a distributed file system 110 comprising multiple “smart” storage units 114. (Mikesell Paragraph [0062]; Figure 1). Data files are broken into blocks and spread or striped across multiple storage devices.

The intelligent distributed file system 110 enables blocks of an individual file to be spread across multiple smart storage units 114.

(Mikesell Paragraph [0062]). The storage devices on which to store the blocks of data are selected based on default parameters or by a system administrator. (Mikesell Paragraph [0089]). However, the block placement does not appear to be controlled based on preferences for the particular file being stored.

Claim 15, by contrast, specifically recites, “wherein the server stores the file on at least one of the first and second storage devices in a global namespace, the selection of the storage location made by the server based on the attributes of the storage devices and storage preferences for the file.” Applicant respectfully submits that Mikesell does not expressly or inherently teach such a system. In Mikesell, data files are broken into blocks, and the blocks are stored in a plurality of different locations, but the locations are selected general settings or selections on not based on preferences for the particular file. Thus, Mikesell does not expressly or inherently teach “wherein the server stores the file on at least one of

the first and second storage devices in a global namespace, the selection of the storage location made by the server based on the attributes of the storage devices and storage preferences for the file.”

Based on the foregoing, Applicant respectfully submits that claim 15, and all claims which depend from claim 15 (claims 16-18), should be allowed.

D. Claim 19

Claim 19 stands rejected as allegedly obvious over Cannon and Mikesell.

Cannon is directed to a storage management system with file aggregation supporting multiple aggregated file counterparts. (Cannon Title). In particular, Cannon discusses the existence of processing overhead associated with a file in a system, and that for smaller files the processing overhead is a predominant factor in accessing the smaller files. (Cannon Col. 1, line 65 through Col. 2, line 7). In order to decrease the effect of the processing overhead associated with each access to a file, Cannon discloses a data storage subsystem 102 that performs services for client stations 106 (*e.g.*, archive and backup) where user files are aggregated into managed files. (Cannon Col. 7, lines 54-67). Cannon discloses tables that track the membership of various user files in a managed file; however, the user files are consistently identified by the name assigned by client stations 106. (Cannon Col. 8, lines 29 through Col. 9, line 11; Table 1 (note “user file name” column); Table 3 (note “user file name” column)).

Claim 19, by contrast, specifically recites, “wherein the first means for executing communicates files to the second means for executing for storage on at least one of the plurality means for storing, wherein the second means for storing to be a network storage device operating in a file structure of the first means for executing; and wherein program executing on the second means executing selects on which of the plurality of means for storing to store the files on a file-by-file basis based on storage characteristic preferences supplied for each file, and wherein each file is stored under a globally unique name in a global namespace of the plurality of means for storing.” Applicant respectfully submits that Cannon and Mikesell do not teach or suggest such a system. In particular, Cannon appears to teach tracking a file throughout the Cannon system by the file name

assigned by the client station. Thus, even if hypothetically the teachings of Mikesell are precisely as the Office action suggests (which Applicant does not admit), Cannon and Mikesell fail teach or suggest "wherein each file is stored under a globally unique name in a global namespace of the plurality of means for storing."


Based on the forgoing, Applicant respectfully submits that claim 19, and all claims which depend from claim 19 (claims 20-25), should be allowed.

V. CONCLUSION

In the course of the foregoing discussions, Applicant may have at times referred to claim limitations in shorthand fashion, or may have focused on a particular claim element. This discussion should not be interpreted to mean that the other limitations can be ignored or dismissed. The claims must be viewed as a whole, and each limitation of the claims must be considered when determining the patentability of the claims. Moreover, it should be understood that there may be other distinctions between the claims and the cited art which have yet to be raised, but which may be raised in the future.

Applicant respectfully requests reconsideration and that a timely Notice of Allowance be issued in this case. It is believed that no extensions of time or fees are required, beyond those that may otherwise be provided for in documents accompanying this paper. However, in the event that additional extensions of time are necessary to allow consideration of this paper, such extensions are hereby petitioned under 37 C.F.R. § 1.136(a), and any fees required (including fees for net addition of claims) are hereby authorized to be charged to Hewlett-Packard Development Company's Deposit Account No. 08-2025.

Respectfully submitted,



HEWLETT-PACKARD COMPANY
Intellectual Property Administration
Legal Dept., M/S 35
P.O. Box 272400
Fort Collins, CO 80527-2400

Mark E. Scott
PTO Reg. No. 43,100
CONLEY ROSE, P.C.
(512) 391-1900 (Phone)
(512) 320-9181 (Fax)
ATTORNEY FOR APPLICANT